
Proceedings Paper:
Mitchinson, B., Sullivan, J.C., Pearson, M.J. et al. (2 more authors) (2013) Perception of simple stimuli using sparse data from a tactile whisker array. In: Biomimetic and Biohybrid Systems. Second International Conference, Living Machines 2013, July 29 – August 2, 2013, London, UK. Lecture Notes in Computer Science, 8064 . Springer International Publishing , pp. 179-190. ISBN 9783642398018

https://doi.org/10.1007/978-3-642-39802-5_16

Perception of simple stimuli using sparse data from a tactile whisker array

Ben Mitchinson[1⋆], J. Charles Sullivan[2], Martin J. Pearson[2], Anthony G. Pipe[2], and Tony J. Prescott[1]

[1] The University Of Sheffield, Sheffield, UK.
[2] Bristol Robotics Laboratory, Bristol, UK.

Abstract. We introduce a new multi-element sensory array built from tactile whiskers and modelled on the mammalian whisker sensory system. The new array adds, over previous designs, an actuated degree of freedom corresponding approximately to the mobility of the mystacial pad of the animal. We also report on its performance in a preliminary test of simultaneous identification and localisation of simple stimuli (spheres and a plane). The sensory processing system uses prior knowledge of the set of possible stimuli to generate percepts of the form and location of extensive stimuli from sparse and highly localised sensory data. Our results suggest that the additional degree of freedom has the potential to offer a benefit to perception accuracy for this type of sensor.

Keywords: Tactile sensing, Whiskers, Perception, Robotics.

1 Introduction

Many small mammals (including rats, mice, and shrews) sport an array of long, touch-sensitive, facial whiskers which they rely on to locate and identify objects and obstacles in their environment [1]. This system has long been a popular model of biology amongst neuroscientists, but it also has characteristics attractive to biomimetic engineers. Most obviously, whiskers operate without the use of any sort of radiation. This means that they can be operated within opaque fluids (e.g. turbid water or smoke-filled air, which would stymy visible-light sensing) and that they are not affected by ambient radiative noise. Less obviously, whiskers report local geometry very directly, so that the well-known ambiguities familiar from, say, visual sensory processing do not arise, and the recovery of geometry is computationally trivial. Thus, artificial sensory whiskers have more recently been built and studied both as individual sensors and as sensory arrays; see Prescott *et al.* [2] for a recent review.

The design of our own multi-whisker robotic platforms has increased in sophistication—see Figure 1 for illustrative examples—alongside our understanding both of the biological whisker sensory system and of the importance of morphological and functional features to an artificial system. The whiskers of

⋆ Corresponding author b.mitchinson@shef.ac.uk.

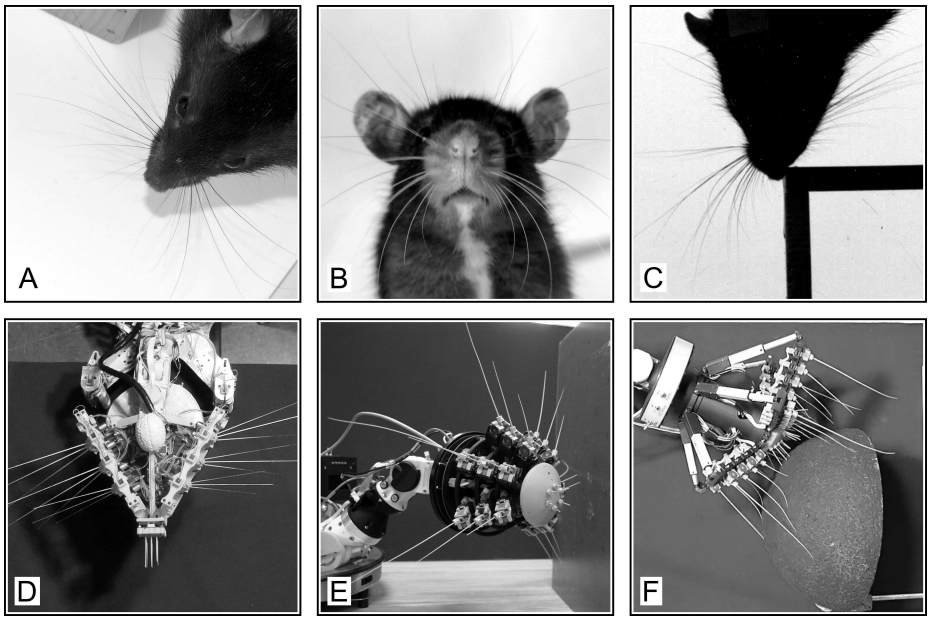

Fig. 1. *Rats and robots.* Features of the mystacial whisker morphology/functionality of rats include bilateral whisker arrays (A), coverage all round the snout (B), and reshaping the mystacial pad to move the bases of the whiskers in response to local geometry (C). These features, respectively, have inspired design features of three generations of whiskered robot: Scratchbot (D), Shrewbot (E) and the BIOTACT G2 Sensor (F).

'Scratchbot' (Figure 1D) were more-or-less sideways-facing, had limited degrees of freedom (DOFs; three per side) [2], and the platform was used to investigate active sensing on small temporal and spatial scales [3]. Scratchbot exhibited 'blind spots', however, owing to the limited coverage of the area around the snout by whiskers, so that unsupervised exploration was impossible. 'Shrewbot'[3] (Figure 1E) added all-around-the-snout whisker sensing, as well as independent actuation of each whisker so that active sensing could be optimised per-whisker [4]. This allowed unsupervised operation on long timescales (for instance, to study tactile SLAM [5]). Using the Shrewbot platform we also demonstrated simultaneous identification and localisation of 3D shapes in a mobile target acquisition task with feedback-controlled interactions [6]. In that study, we used model-based percepts to drive effective animal-like prey capture behaviour from sparse, localised, sensory data. Alongside Shrewbot, we developed a system—the 'BIOTACT G1 Sensor' (not shown in the Figure)—that mounted a similar array of whiskers on a robot arm rather than a mobile platform to facilitate investigations of precisely controlled interactions [4]. G1 has been used to demonstrate identification/localisation of both textured surfaces [7] and 2D shapes [8, 9].

[3] Though, as the name implies, the design of Shrewbot is actually based on a type of shrew, the functional and morphological features are much the same in rats.

A shortcoming of these systems that became apparent through experience is the difficulty in bringing many whiskers into contact with convex local geometry (for example, small objects in front of the snout). This is because the bases of the whiskers are arranged in a regular fashion spread along—essentially—the 'length' of the robot, and more rearward whiskers cannot cross over more forward whiskers. Observation of rats exploring such challenging geometries (for example, see Figure 1C) suggests that quite different arrangements of the whisker bases are available to the animal depending on what is being sensed, through control over the shape of the 'mystacial pad' (something like the 'cheek'). A recent, very detailed, study of rat whisker and facial morphology highlighted the centrality of such morphological details to the available sensor-environment interactions with objects of different curvature [10]. In the current study, then, we introduce the next generation BIOTACT Sensor: 'G2' (Figure 1F) adds an additional DOF to each whisker row—control over the angle of the bearer ('cheek') which carries the whisker modules—intended to offer more flexible sensor-environment interactions. We also report the results of a task involving identification/localisation of 3D shapes with four stimulus classes and pre-programmed interactions, confirming and extending our previous studies using Shrewbot and G1, and providing a first assessment of the usefulness of the additional DOF.

2 Methods

2.1 Data collection

Hardware G2 is shown in Figure 2 with 18 whisker modules fitted (3 whiskers in each of six rows distributed regularly around the 'snout'). The whisker modules are identical to those used in G1 [4]. Briefly, the module integrates a composite tapered whisker mounted on a rotating shaft, motor and motor controller to cause whisker 'protraction', and Hall effect sensors measuring protraction angle and 2D whisker deflection. Each whisker can be rotated around its base through approximately 100° (Figure 2A/B). Each whisker module is mounted on one of six cheek members, pivoted near the tip of the snout, which can be swung backwards and forwards through approximately 30° (Figure 2B/C) by a linear actuator (Firgelli L12, `firgelli.com`). The whiskers vary in length along each row, from 80mm (at the front) to 165mm (at the back). The G2 sensor is mounted on a 7 DOF robotic arm so that it can be positioned and oriented as required in space. Note that the sensor also sports an array of small whiskers at the front in the centre, visible in the Figures, that are not used in this study.

Experimental protocol To perform experimental trials, the sensor is brought by the robot arm into juxtaposition with one of four stimuli, S1-4, so that the sensor 'faces' the stimulus (Figure 2D). S1-4 are, respectively, spheres of diameter 200, 300 and 600mm, and a flat plane. The location of the sensor relative to the stimulus is controlled in one lateral dimension to ten locations, L1-10, 40mm apart. At each position, the sensor conducts thirty 'whisks'. Each whisk takes one second, and involves protracting all the whiskers forward until they either reach

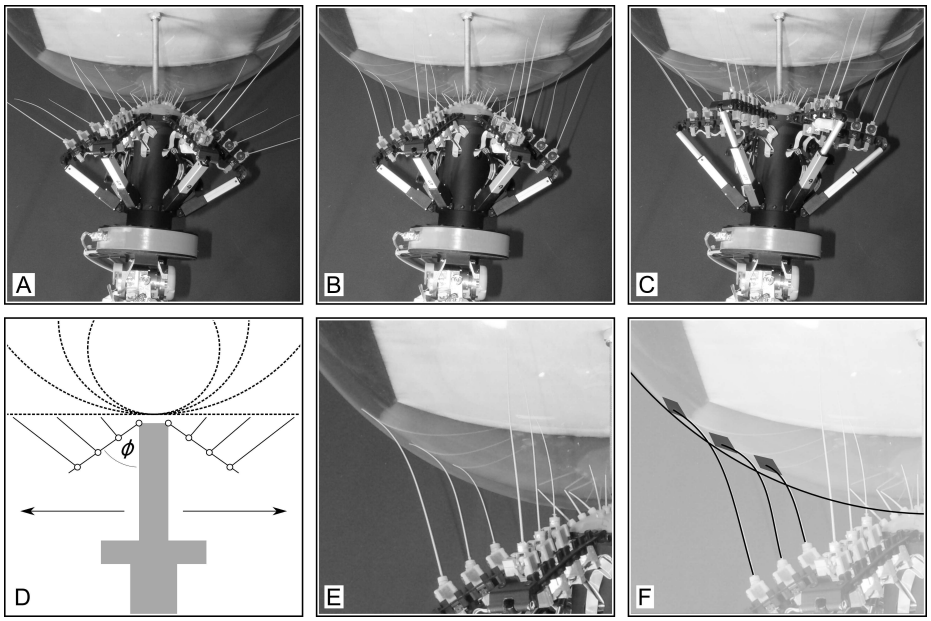

Fig. 2. *G2 Sensor.* Top-down views of G2 contacting stimulus S3 (600mm sphere), with cheeks and whiskers retracted (A), whiskers protracted (B), and cheeks and whiskers protracted (C). *Experimental protocol.* (D) Top-down view of experimental protocol showing key functional geometry of G2 (solid lines; cheeks and whiskers) as well as the four stimuli (dashed lines; three spheres and one plane, S1-4). The variable relative location of sensor and stimulus in one dimension (L1-10) is indicated by arrows; the cheek protraction angle—denoted ϕ—is also indicated. *Sensory perception.* (E) Close-up of whisker-stimulus interaction from panel B. (F) Same image with stimulus and three whiskers highlighted (solid lines) and three 'surfels' indicated (dark grey patches).

their maximum protraction angle (after 700ms) or contact a surface; in either event, the individual whisker is retracted. This 'early' cessation of protraction in the case of surface contact is intended to implement a 'minimal impingement' control strategy such as has been proposed as a component of the whisker control strategy used by rats [11]. The cheek angles (ϕ, Figure 2D) are controlled according to one of three strategies, C1-3 (described below), in separate realisations of the protocol. Thus, this study involves a total of $4 \times 10 \times 30 \times 3 = 3600$ whisks, which we hereafter label as trials. We index trials with the variable k, and denote the stimulus for trial k as s_k.

Contact signals At each time sample (200Hz), each whisker reports a 2D deflection vector, $\boldsymbol{d}_w$(t). From this, we derive a 'contact belief' signal, $b_w(t)$, as

$$b_w(t) = \begin{cases} 0, v_w(t) < 0 \\ 1, v_w(t) > 1 \\ v_w(t), \text{otherwise} \end{cases}, \quad v_w(t) = ||\boldsymbol{d}_w(t)||/\eta_w - 1. \tag{1}$$

The parameter η_w is a measure of the sensor noise during non-contacting whisking, and is chosen separately for each length of whisker (of which there are three) based on data from test trials where no stimulus was present: η_w is set by eye to the peak value of $||\boldsymbol{d}_w(t)||$ seen in such test trials plus about 50%. Thus, values of b_w of 0 and 1 are intended to indicate, approximately and respectively, 'certainty of not contact' and 'certainty of contact'.

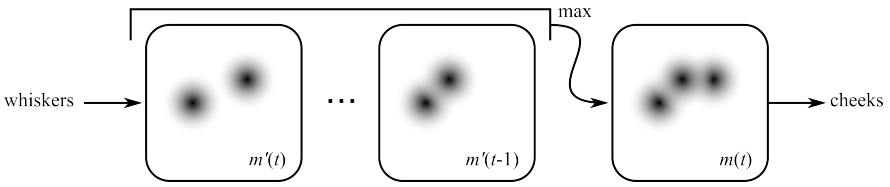

Fig. 3. *Egocentric processing.* 'Whisker-centric' contact signals are mapped into a 3D egocentric map $m'(t)$ through a Gaussian spatial filter. An element-wise maximum over one second's worth of this signal stream forms the (with memory) contact map $m(t)$, which drives the cheeks when the C1 control strategy is in use.

Egocentric processing Spatial processing is performed using egocentric maps, each of which is maintained as a 3D array of voxels (cubes of side 10mm). Using the known geometry of the sensor and the instantaneous cheek and whisker angles, the locations of the whisker tips in 3D egocentric space are computed. The contact belief signals for all whiskers, $\{b_w(t)\}$, are then transformed into a map $m'(t)$ at these tip locations through a Gaussian spatial filter of width 25mm, to form an egocentric map of instantaneous contact belief (see Figure 3). An element-wise maximum operation over all samples of $m'(t)$ from the last one second is then applied, to produce a single final map of contact belief, $m(t)$, that has a one second memory (Figure 3). This map $m(t)$ forms the input to a transform, described below, that controls the cheeks for the C1 strategy.

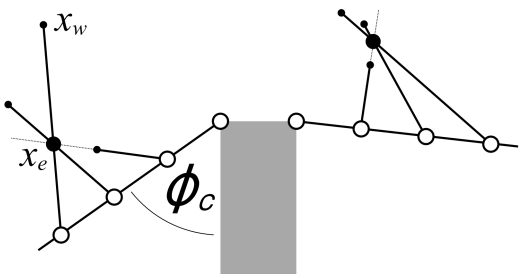

Fig. 4. *Strategy C1.* (Left) Example element of egocentric map at $\boldsymbol{x}_e$ (large black dot). For cheek angle shown, ϕ_c, whiskers are adjusted so tips (small black dots) are as close to $\boldsymbol{x}_e$ as possible. Tip of wth whisker in that configuration is denoted $\boldsymbol{x}_w(\phi_c)$. (Right) A second example, at a different location.

Cheek control strategies The C1 control strategy chooses a protraction angle independently for each cheek such that the whiskers mounted on it are afforded

the opportunity, as far as possible, to bring their tips to where contact has recently occurred (i.e. to regions of $m(t)$ that are active). The strategy is written

$$\phi_c = \frac{\sum_e K_c \star U_c \star m(t)}{\sum_e U_c \star m(t)}, \tag{2}$$

where ϕ_c is the angle of the cth cheek, K_c is a 3D array of proposed cheek angles for the cth cheek equal in size to $m(t)$, U_c is a weight array of the same size, $\star$ represents element-wise multiplication and $\sum_e$ is the sum over all elements of the 3D array. Thus, ϕ_c is given by an average weighted by $U_c \star m(t)$ over the angles proposed by each element of K_c. The arrays K_c and U_c are generated as follows (see also Figure 4). To compute the eth element of both arrays, we start with the location in space represented by that element, denoted $\boldsymbol{x}_e$. For any value of ϕ_c, the protraction angle of each whisker can be adjusted so as to bring its tip as close as possible to $\boldsymbol{x}_e$. Its tip position in this configuration is denoted $\boldsymbol{x}_w(\phi_c)$. The 'benefit' of a proposed cheek angle for this whisker for this element is then calculated as $u_{e,w}(\phi_c) = \exp(-||\boldsymbol{x}_e - \boldsymbol{x}_w(\phi_c)||/r)$ where $r = 75$mm. That is, if the tip of the whisker can reach near to the element's location at some whisker angle, given the proposed cheek angle, a large benefit results; otherwise, a small benefit results. Finally, the preferred cheek angle ϕ_c for this element e of the map is simply that which maximises the average benefit $u_{e,w}(\phi_c)$ across the whiskers w of that cheek. This angle is entered into the corresponding element of K_c, and the average benefit at that value is entered into U_c. To summarize, activity in a region of the map $m(t)$ will tend to drive a cheek so that the tips of the whiskers mounted on it can reach that region (encoded in K_c), but only for regions of the map for which that can actually be achieved (encoded in U_c). Thus, strategy C1 is intended to favour cheek angles that permit whisker-stimulus contact at *some* whisker protraction angle, so as to implement a 'maximal contact' control strategy, proposed by Mitchinson *et al.* [11] as a component of the strategy used by rats. Strategies C2 and C3 set all cheek protraction angles to a fixed value, as if the cheeks were not actuated; the three strategies are summarised below.

- **C1** Cheek angles controlled individually and automatically based on past sensory data to favour contact at the whisker tips (Equation 2).
- **C2** All cheek angles fixed at fully protracted ($\phi_c = 85° \; \forall \; c$).
- **C3** All cheek angles fixed at halfway protracted ($\phi_c = 70° \; \forall \; c$).

2.2 Analysis

Feature extraction The data streams from each whisker are initially treated separately. Within each trial, for the wth whisker, we identify the time of maximum protraction angle, and sample $\boldsymbol{d}_w$ and b_w at that time (we expect, then, to have taken these samples when the whisker was undergoing the most strong deflection). For each whisker a 'surfel'—an oriented patch of surface—is computed, as follows (some example surfels are indicated in Figure 2F). First, the amount of physical deflection of the whisker tip is estimated according to $\boldsymbol{v}_w = q_{\boldsymbol{v}} \boldsymbol{d}_w$

where $q_{\boldsymbol{v}}$ is an (unknown) gain between deflection sensor data and physical deflection in millimetres. The undeflected location of the whisker tip is calculated according to the known geometry and the current angles of cheek and whisker, and $\boldsymbol{v}_w$ is added to it to give the estimated location of the whisker tip at peak protraction/deflection, $\boldsymbol{x}_w$. Next, a 'surface parallel' vector, $\boldsymbol{p}_w$, is estimated by finding a unit vector perpendicular to both $\boldsymbol{v}_w$ and the undeflected whisker shaft (note that the surface normal cannot be uniquely identified from the report of a single whisker at a single time sample). Finally, the surfel $L_w = \{\boldsymbol{x}_w, \boldsymbol{p}_w, b_w\}$ consists of an estimate of a point on the surface, $\boldsymbol{x}_w$, an estimate of a unit vector parallel to the surface, $\boldsymbol{p}_w$, and a measure of the belief that contact occurred, b_w. This is, of course, a simplified model of whisker deflection but is expected to be reasonably accurate for the small deflections and contact at or near the tip that are encountered in this experiment.

Candidate percepts For the kth trial, the complete set of sensory data extracted can be written as $\mathbf{L}_k = \{L_w\}$, one surfel for each whisker. Four percept models, denoted $\hat{\mathrm{S}}$1-4, are fitted to these data. $\hat{\mathrm{S}}i$ corresponds to Si—that is, $\hat{\mathrm{S}}$1-3 are models of spheres with radii of 200/300/600mm, whilst $\hat{\mathrm{S}}$4 is a model of a plane. The parameters of $\hat{\mathrm{S}}$1-3 are the location relative to the tip of the snout of the centre of the sphere; the parameters of $\hat{\mathrm{S}}$4 are the distance ahead of the tip of the snout and the angle the plane normal makes with the midline of the snout. The parameters of each model $\hat{\mathrm{S}}i$ are optimised (`fminsearch`, Mathworks MATLAB$^{\mathrm{TM}}$) against $\mathbf{L}_k$ for each trial using a cost function

$$J(\hat{\mathrm{S}}i, \mathbf{L}_k) = \sum_w \Big(J_1(\hat{\mathrm{S}}i, L_w) + J_2(\hat{\mathrm{S}}i, L_w) + q_{\boldsymbol{p}} J_3(\hat{\mathrm{S}}i, L_w) \Big) \tag{3}$$

$$J_1(\hat{\mathrm{S}}i, L_w) = b_w |F(\hat{\mathrm{S}}i, \boldsymbol{x}_w)| \tag{4}$$

$$J_2(\hat{\mathrm{S}}i, L_w) = (1 - b_w) \max(F(\hat{\mathrm{S}}i, \boldsymbol{x}_w), 0) \tag{5}$$

$$J_3(\hat{\mathrm{S}}i, L_w) = b_w \left| < G(\hat{\mathrm{S}}i, \boldsymbol{x}_w), \boldsymbol{p}_w > \right|. \tag{6}$$

In the above, both $F(.)$ and $G(.)$ are functions that involve computing the point on the surface of the model $\hat{\mathrm{S}}i$ nearest the surfel location $\boldsymbol{x}_w$, denoted $\boldsymbol{x}'_w$. $F(.)$ returns the signed distance that $\boldsymbol{x}_w$ is inside the model surface (i.e. the distance from $\boldsymbol{x}_w$ to $\boldsymbol{x}'_w$, when $\boldsymbol{x}_w$ is inside $\hat{\mathrm{S}}i$, and the negative of this value, otherwise). $G(.)$ returns the normal to the model surface at $\boldsymbol{x}'_w$. We do not explicitly describe these functions since they are simple geometric computations specific to the form of each percept model (sphere or plane, here). The operator $<,>$ indicates the scalar product. Thus, the three cost function components penalise, respectively, surfels that are distant from the model surface and have high contact belief, surfels that are *inside* the model surface and have low contact belief, and surfels that have high contact belief and are not aligned with the model surface. The overall cost, J, is a weighted sum of the three components, where $q_{\boldsymbol{p}}$ is an (unknown) weight parameter for the alignment cost component.

Unitary percept Finally, a unitary percept for the kth trial, $\hat{s}_k$, is generated simply by selecting the candidate percept $\hat{\mathrm{S}}i$ with the lowest weighted cost.

That is,

$$\hat{s}_k = \hat{\mathrm{S}}i,\ i = \underset{i}{\operatorname{argmin}} \Big(q_i J(\hat{\mathrm{S}}i, \mathbf{L}_k) \Big), \tag{7}$$

where $\{q_i\}$ are a set of (unknown) weights. These weights are required, for example, because smaller stimuli may elicit lower cost values simply by eliciting a lower number of whisker contacts. Note that we do not need to consider an alternative percept 'no stimulus present' since our experiment does not permit that case (cf. a related investigation where absence of stimulus is a valid case [6]). We can then compute the 'correctness' of the identification,

$$C_k = \begin{cases} 1, \hat{s}_k = s_k \\ 0, \hat{s}_k \neq s_k \end{cases} . \tag{8}$$

3 Results

3.1 Data collection

In these experiments the whiskers were used to palpate the stimuli, as described in Methods. Therefore, sensory data (non-zero values of contact belief) are available for a short period (some tens of milliseconds) within each whisk/trial (one second), at around the time of maximum whisker protraction. The one second memory in $m(t)$ is necessary so that these transient data can govern the cheek angles during the next whisker-stimulus interaction, bridging the intervening hiatus in data availability (incidentally, the cheek actuators take around the same period of time—one second—to respond to a commanded angle). As a consequence, within each 30 trial set, the cheek angles during trial k are driven by sensory data collected during trial $k-1$. Since within these sets the stimulus and sensor are not moved, this delay—that would also be encountered by a whisking animal—need not be accounted for by the control algorithm. Such a long delay does permit that a large set of controllers would generate oscillation; with the controller described, however, we generally saw instead rapid (one trial or so) convergence to a fixed cheek position, though we did not assess this formally.

3.2 Parameter optimisation

Table 1. Analysis parameter values chosen by independent optimisation for each condition (C1, C2, C3) and by global optimisation across all conditions (C1-3). $q_1 = 1$ is assumed in all conditions.

Condition	$q_{\boldsymbol{v}}$	$q_{\boldsymbol{p}}$	$[q_2\ q_3\ q_4]$
C1	250	10	[0.75 0.60 0.60]
C2	150	2	[0.70 0.50 0.45]
C3	150	10	[0.85 0.60 0.70]
C1-3	250	10	[0.95 0.90 0.90]

The analysis described above has six unknown parameters: $q_{\boldsymbol{v}}$, $q_{\boldsymbol{p}}$, and $\{q_i\}$, $i \in [1, 4]$. In practice, the percept weights can be normalised, so we set $q_1 = 1$, leaving just five unknowns. $q_{\boldsymbol{v}}$ could in theory be measured, but it is difficult to do so accurately in practice, so we used optimisation to select all five. To make fair comparison across the conditions C1-3, we did this separately for each condition. Thus, we split the 3600 trials collected into training and test sets for each condition. Specifically, of each set of 30 whisks/trials, we discarded the first 2 (control loop settling), used the next 8 for training, and the remaining 20 for testing. We then maximised the total number of 'correct' training trials ($\sum_k C_k$) against the analysis parameters by exhaustive search. This *per-condition* parameter choice included fairly similar values for each condition, as shown in Table 1. We also computed a *global* set of analysis parameters that were optimised across all conditions (C1-3). The remainder of this section reports the results of using these parameter sets to process the 2400 test trials, using the *per-condition* parameter set except where stated.

3.3 Performance

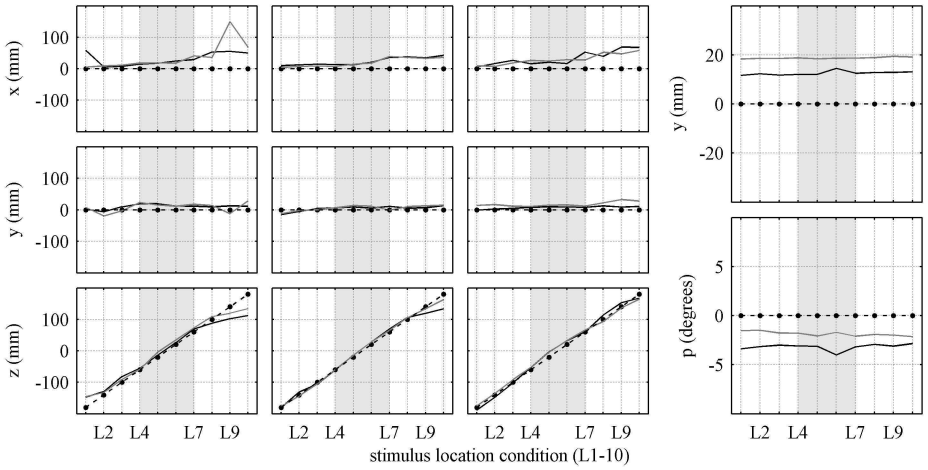

Fig. 5. Columns 1-4 correspond to S1-4 (also, $\hat{\text{S}}$1-4, respectively, since correct identification is assumed, here). Columns 1-3 graph normalised location of sphere models $\hat{\text{S}}$1-3 (ground truth is dashed line, estimates due to C1/2 are solid black/grey lines; estimate due to C3 was qualitatively similar and is omitted for clarity), averaged across 20 test trials. Column 4 shows location and orientation of plane model $\hat{\text{S}}$4 (averaged across whisks, legend as for columns 1-3). Shaded region is L4-7 (see text).

We performed two analyses of the test data. In the first, we provide identification information—that is, we consider only model $\hat{\text{S}}i$ for trials with stimulus Si—and review performance in localisation of the stimuli. Results of this

analysis are graphed in Figure 5, averaged across 20 test trials. Breaking down these results by cheek control strategy, mean localisation error averaged over all sphere trials was 37/38/46mm (for C1/2/3, respectively), over all plane trials was 13/19/20mm and 3/2/2°. The graphs show that localisation is poorer at larger distances from smaller spherical stimuli—as expected, since fewer whiskers contact the stimulus under these conditions. In an 'inner region' L4-7, where multiple contacts are most likely, mean sphere localisation error was 29/31/35mm (location did not, as expected, affect performance in locating the plane stimulus). Additionally, the graphs suggest some systematic sources of error—the standard deviation of localisation error across the 20 trials contributing to each point in the plots is not shown, for clarity, but was smaller than the mean error at the majority of parameter points—we return to this in Section 4. Using the *global* parameter set, sphere localisation error was 37/42/43mm and plane localisation error was 13/7/8mm and 3/2/3° (C1/2/3, respectively).

In the second analysis, the percept $\hat{s}_k$ with the lowest weighted cost was chosen for each trial and compared with the true stimulus identity s_k, as described in the equations above. Identification was correct in 92/90/93% of trials (for C1/2/3, respectively). Focusing only on L4-7, these percentages are 99/97/85%. It is interesting that the performance of C3 is actually substantially worse in the inner region than overall—in an 'outer region' (L1-3 and L8-10), identification was correct in 87/85/99% of trials. Examining the confusion matrices in the C3 case for the inner region indicates that 47 of 48 mis-identifications made were mistaking S2 for S1 (i.e. perceiving a 200mm sphere when a 300mm sphere is present), a matter we will return to in Section 4. Using the *global* parameter set, identification was correct in a slightly reduced 88/89/89% of trials.

4 Discussion

We have presented a new multi-element tactile sensory array with an additional DOF (per row of whiskers) over previous sensors, as well as a preliminary quantitative analysis of identification and localisation of four simple stimuli using a model-based technique. Overall, our results confirm and extend those of a previous—smaller and less controlled—study using two stimuli and a mobile platform [6]. That is, this type of sensor and analysis can discriminate amongst such a set of objects with correctness around 90% and, given correct identification, simultaneously locate them with precision on the order of a few tens of millimetres. As illustrated, model-based perception is particularly well suited to whisker sensory systems since it generates complete percepts of extensive stimuli from the highly localised sensory data available from the whisker array.

We also made a first attempt at deriving a control strategy for the additional DOF that favours optimal sensing. Such a strategy could be derived by experiment, but the procedure would be onerous (owing to the large space of possible stimuli). Here, we hand-crafted a controller (C1) intended to favour contact at the tips of the whiskers. Compared with strategy C2 (cheeks fully protracted), C1 seemed to perform marginally better in both localisation and identification

analyses. Compared with strategy C3 (cheeks half protracted), the picture was more complex. C1 performed consistently (and more markedly) better at localisation, and also at identification for the 'inner region' (L4-7), where the sensor was most directly opposite the stimulus (99% correct for C1 versus 85% for C3). However, in the 'outer region' (L1-3, L8-10), where fewer whiskers would have been able to contact the stimulus in general, the picture for identification was reversed, with C3 (99% correct) being markedly superior to C1 (87% correct). Interestingly, the great majority of the mistakes made in the C3/inner case (47 of 48) were confusing the second smallest stimulus, S2, for the smallest, S1. Most of the mistakes made in the C1/outer case (55 of 61) were mistaking the smallest sphere, S1, for one of the other stimuli. These results confirm the general point that smaller (or, perhaps, more convex) stimuli are the more challenging to sense for a system with this general morphology.

A confounding systematic source of error is visible in the plots of Figure 5, both in the trend in the x location of spheres and in the fixed offset in the y location of the plane. The value of the fixed offset was affected by the choice of analysis parameters, underlining one probable source of error: imprecision in the whisker deflection model. Another probable source is error in calibration of hardware geometry, of the robot arm, whiskers, or stimuli. The relative error of the localisation measurements made by the system was generally lower in magnitude than the absolute systematic error, so that it is difficult to draw reliable conclusions regarding the relative performance of the different controllers. Overall, then, these results seem too complex to draw strong conclusions about the benefit to sensing offered by the online cheek control strategy tested here (C1). However, they do indicate unambiguously that a *post hoc* control strategy exists that would substantially outperform either of the 'non-actuated' strategies C2 or C3 in identification. That is, to use C1 in the inner region (99% correct) and C3 in the outer region (99% correct). This strategy markedly out-performs all other strategies tested here (99% correct, versus 92/90/93% correct for C1/2/3).

One line of future work, then, will be to seek a controller that can realise this *post hoc* advantage at run-time. In addition, the conditions under which we would expect to see most strongly the advantage of actuated cheeks remain to be tested: these are when stimuli can be presented not only ahead of, but also to one side of, the sensor. An experiment to test under these conditions is more difficult to design, since some sensor/stimulus/cheek combinations will generate collisions between the sensor hardware and the stimulus. Nonetheless, such tests will be necessary if the full gamut of conditions experienced by the exploring animal/mobile-robot is to be addressed. Figure 5 highlights that the location of the sensor relative to the stimulus (L1-10) has an impact on localisation performance. Lepora *et al.* have shown that active control over location using sensory feedback at run-time can improve sensing performance under such conditions [12] (this issue); incorporating such feedback should improve localisation performance of this system, also. Whilst identification performance using each individual strategy did appear to be location-dependent, the very high performance (99% correct) available from the *post hoc* C1/C3 strategy means that

it is impossible to gauge whether active sensing of this type would also aid identification (though it seems intuitively likely). Accordingly, any future tests of identification using active sensing will need to employ more finely-graded stimuli. Finally, the use of an improved whisker deflection model and the addition of adaptive calibration may be able to eliminate systematic errors. We expect inclusion of these features to permit substantially more discerning identification and more precise localisation using a whisker array of this type.

Acknowledgements The authors thank the technical support engineers at Bristol Robotics Laboratory and the anonymous reviewers. This work was supported by the FP7 grants BIOTACT (ICT-215910) and EFAA (ICT-270490).

References

1. Brecht, M., Preilowski, B., Merzenich, M.M.: Functional architecture of the mystacial vibrissae. Behavioural brain research **84**(1) (1997) 81–97
2. Prescott, T.J., Pearson, M.J., Mitchinson, B., Sullivan, J.C.W., Pipe, A.G.: Whisking with robots. Robotics & Automation Magazine, IEEE **16**(3) (2009) 42–50
3. Pearson, M., Mitchinson, B., Welsby, J., Pipe, T., Prescott, T.: Scratchbot: Active tactile sensing in a whiskered mobile robot. Animals to Animats 11 (2010) 93–103
4. Pearson, M.J., Mitchinson, B., Sullivan, J.C., Pipe, A.G., Prescott, T.J.: Biomimetic vibrissal sensing for robots. Philosophical Transactions of the Royal Society B: Biological Sciences **366**(1581) (2011) 3085–3096
5. Pearson, M.J., Fox, C., Sullivan, J.C., Prescott, T.J., Pipe, T., Mitchinson, B.: Simultaneous localisation and mapping on a multi-degree of freedom biomimetic whiskered robot. To appear, ICRA 2013.
6. Mitchinson, B., Pearson, M.J., Pipe, A.G., Prescott, T.J.: Biomimetic tactile target acquisition, tracking and capture. Robotics and Autonomous Systems, *in press*.
7. Sullivan, J., Mitchinson, B., Pearson, M.J., Evans, M., Lepora, N.F., Fox, C.W., Melhuish, C., Prescott, T.J.: Tactile discrimination using active whisker sensors. Sensors Journal, IEEE **12**(2) (2012) 350–362
8. Lepora, N., Fox, C., Evans, M., Mitchinson, B., Motiwala, A., Sullivan, J., Pearson, M., Welsby, J., Pipe, T., Gurney, K., et al.: A general classifier of whisker data using stationary naive bayes: application to biotact robots. Towards Autonomous Robotic Systems (2011) 13–23
9. Lepora, N.F., Sullivan, J.C., Mitchinson, B., Pearson, M., Gurney, K., Prescott, T.J.: Brain-inspired bayesian perception for biomimetic robot touch. In: IEEE International Conference on Robotics and Automation (ICRA), (2012) 5111–5116
10. Towal, R.B., Quist, B.W., Gopal, V., Solomon, J.H., Hartmann, M.J.: The morphology of the rat vibrissal array: a model for quantifying spatiotemporal patterns of whisker-object contact. PLoS Computational Biology **7**(4) (2011) e1001120
11. Mitchinson, B., Martin, C.J., Grant, R.A., Prescott, T.J.: Feedback control in active sensing: rat exploratory whisking is modulated by environmental contact. Proceedings Royal Society B: Biological Sciences **274**(1613) (2007) 1035–1041
12. Lepora, N.F., Martinez-Hernandez, U., Prescott, T.J.: A SOLID case for active bayesian perception in robot touch. In: Proceedings of Living Machines. (2013)